REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
12/31/2010	Annual	2008.1.1. -- 2010.12.31

4. TITLE AND SUBTITLE
Towards an Analytic Foundation for Network Architecture

5a. CONTRACT NUMBER
N00014-07-1-0864

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Chiang, Mung

5d. PROJECT NUMBER
10PR00613-00

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Princeton University
Engineering Quadrangle
Olden Street
Princeton, NJ 08544

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Office of Naval Research
875 North Randolph Street
Arlington, VA 22203-1995

10. SPONSOR/MONITOR'S ACRONYM(S)
ONR

11. SPONSORING/MONITORING AGENCY REPORT NUMBER

12. DISTRIBUTION AVAILABILITY STATEMENT
Approved for public release; distribution is unlimited

13. SUPPLEMENTARY NOTES
N/A

14. ABSTRACT
In this project, we develop the analytic tools of stochastic optimization for wireless network design and apply them to operations of tactical networks. We obtain fundamental results on performance and stability of network resource allocation under flow, packet, and channel dynamics. Distributed and optimal scheduling and power control, in particular, have been developed. We construct a first-principled approach to design network control protocol stack.

15. SUBJECT TERMS
Wireless networks, communications, control, optimization

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
			None		Mung Chiang
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPONE NUMBER (*Include area code)*
					609-258-5071

ONR YIP: Towards an Analytic Foundation for Network Architecture

Annual Report, December 2010

1 Project Participants at Princeton

1.1 Principal Investigators

Mung Chiang, Princeton University

1.2 Post-Doctoral Researchers

Dr. Hazer Inaltekin
Dr. Rui Zhang-Shen

1.3 Graduate Students

Wenjie (Joe) Jiang
Ying Li

1.4 Collaborators and Organizational Partners

Collaboration with Qualcomm, Motorola, and Microsoft.

2 Activities and Findings

2.1 Major Research Activities

The following results have been presented in multiple plenary and keynote speeches given by the PI, including those at IEEE INFOCOM, IEEE GLOBECOM, IEEE WiOpt, KU Leuven Simon Steven Lecture, and MIT LIDS Student Conference.

First is on scheduling algorithms in wireless networks. In an interference environment like wireless cellular or ad hoc networks, scheduling controls which link can transmit at each time slot and has been extensively studied since the 1960s. After the 1992 paper on maximum weight scheduling by Tassiulas and Ephremides, the research community has been developing simpler and more distributed scheduling algorithms that can still perform well in terms of throughput and delay. In an ACM Mobihoc paper 2008, we developed the first unifying framework of throughput-delay-complexity 3 dimensional tradeoff. Each scheduling algorithm is represented as a point in the 3D tradeoff space, and achievable tradeoff curves are extended through various parameterizations. This provides the first systematic way to compare a variety of scheduling algorithms in a fair way. Then in an ACM Mobihoc 2009 paper, we further developed the technique of heavy traffic approximation for wireless scheduling's delay performance characterization. By proving the so called state space collapse property, we can reduce the dimensionality of the problem from L links to 1 "representative" link. We showed that under heavy load and communication overhead in the scheduling algorithms, delay increases exponentially as the network grows.

Furthermore, in collaboration with Microsoft Research in UK, we were among the three teams that developed in 2008-2009 the first adaptive CSMA algorithm that can approach utility optimality arbitrarily tightly without using any explicit message passing. Each node can learn the interference environment purely based on the observed service rate in the past. Proving this result required substantial innovation in the mathematics of stochastic approximation theory and distributed optimization algorithm. This result also lead to new discoveries on transient behaviors of adaptive CSMA, such as the tradeoff between long-term efficiency and short-term fairness among the interferring links.

Then in a paper to appear to appear in IEEE INFOCOM 2011, together with my former postdoc (now faculty at KAIST Korea) Yung Yi and collaborator Ed Knightly at Rice, we reported the first ever *experimental implementation* of optimal CSMA. We successfully implemented the theory inspired design onto conventional 802.11 drivers, offering a feasible path towards deployment. We also demonstrated the predictive power of theory, while discovering the gaps between theory and practice. This further leads to a redesign by our team that bridged these gaps successfully.

Second line of work is on the practically important question of implementable and sub-optimal algorithms. Protocol components in the current network architecture are often designed to attain certain optimality goals, with the hope that, when these optimal components work together, the overall network performance will also be optimized. However, for many network settings, due to either the scale of the network; the constraint on the response time of the algorithms; or the inherent non-convexity in the system, such optimal solutions can be difficult to attain. In this project, we explore architectural choices that are robust to sub-optimality in each individual component. We argue that there is a need to shift our attention from optimal but complicated solutions, to easily implementable designs that are suboptimal but still possess good performance bounds. In the first year of the project, we studied the rate-allocation component of the network architecture, and investigated the following questions related to suboptimal components. (1) We investigated the robustness of the network architecture by studying how much sub-optimality the rate-allocation component can exhibit while the overall network architecture can still achieve satisfactory user-level performance. (2) We investigated how to tradeoff suboptimal rate-allocation with other performance measures, e.g., throughput and link utilization.

Our findings in a paper to appear in IEEE/ACM Transactions on Networking demonstrate that it is possible to design an overall network architecture that is robust to suboptimal components. In particular, we show that even when the transport layer only computes suboptimal rate allocation, under suitable conditions the system can still achieve good user-level performance (in terms of achieving the largest connection-level stability region). Specifically, when the ratio of the utility gap (caused by a suboptimal rate allocation algorithm) to the maximum utility approaches zero as queue length tends to infinity, the maximum connection-level stability region can be retained. When the utility gap is in proportion to the maximum utility, only a reduced stability region can be achieved, in which case we provide a lower bound for the achievable stability regions. Not only that these results demonstrate how to characterize and design network architectures that are robust to suboptimal (but potentially simpler and easier-to-implement) rate-control, they also allow the network designer to intentionally under-optimize a given design objective, with the goal to improve other performance measures of the network.

Third main line of work has been on wireline backbone network architecture, especially on the gap between content distribution and network resource allocation, what we refer to as the "content-pipe" divide. Continuing with our initial work last year, we have in an ACM Sigmetrics 2009 paper a complete study on the interaction between Content Delivery Networks and Internet Service Provider. We show that the current practice of separation between server selection by CDN and traffic engineering by ISP can reach a Nash equilibrium, but a suboptimal one. Sharing information between the two entities about their individual optimization may not always improve efficiency either. In fact, CDN having access to information about network topology and routes from the ISP may actually hurt CDN performance. However, in a clean-slate design where sharing control is possible, we show that by using a Nash bargaining solution that ensure both Pareto-efficiency and fairness, we can redraw the architectural boundary between overlay server selection by CDN and underlay traffic management by ISP so that an efficient Nash equilibirum can be achieved with limited message passing.

Fourth is a theoretical foundation for adaptive network virtualization. Network virtualization has emerged as a powerful way to allow multiple network architectures, each customized to a particular application or user community, to run on a common substrate. Each virtual network could run its own protocols to make efficient use of its share of the underlying resources. However, running multiple virtual networks in parallel raises several key questions: Can each of the network architectures be designed completely independently, without regard for the other virtual networks that would run in parallel with them? How should the shared resources, such as link bandwidth, be divided between the multiple virtual networks, and on what timescale? How does the resulting system compare to a single monolithic design that tries to meet the needs of the multiple applications?

To answer these questions, our work draws on recent advances in using optimization theory to "derive" network protocols. The key insight underlying our work is that *primal decomposition*—a common solution technique in optimization theory—essentially corresponds to network virtualization, with the ability to dynamically vary the share of the resources allocated to each virtual network. The beauty of primal decomposition is that the two child problems can now be solved independently to generate two separate protocol designs, each customized to the corresponding traffic class, while ensuring that the resulting solutions collectively maximize the original joint objective.

Fifth is the theoretical foundation for content delivery over the Internet. Despite the widespread use of P2P technologies for video streaming in the Internet today, the fundamental limit of the highest achievable rate through any P2P method (tree, mesh, push, or pull-based peering) remains unknown until several recent papers. In these papers, we establish a taxonomy of 16 variants of the wireline P2P streaming capacity problem, and develop a tree-based peering construction that is proved to achieve the capacity in 8 of the 12 cases where capacity is unknown before. During the process, a suite of combinatorial graph-theoretic algorithms have been developed. Furthermore, practical scheduling algorithms for wireless P2P systems are developed.

2.2 Training and Development

Post-doc Rui Zhang-Shen will join the networking group at Google. Graduate student Joe Jiang will intern at Thomson Research in Paris in the fall semester on peer-to-peer protocols, and has been working with Gary Chan from Hong Kong University of Science and Technology on applying his ideas to an operational peer-to-peer IPTV system. Graduate student Ying Li graduated last year and joined Samsung wireless standardization group in Texas.

2.3 Outreach Activities

Under-represented groups: At Princeton, the project has involved a female graduate student in engineering: Ying Li.

Professional outreach: The subject of wireless scheduling was recently given as a keynote speech at IEEE WiOpt Conference in Seoul, Korea, by Chiang. Chiang has given numerous talks on architectures for future Internet and wireless networks, to outline the intellectual challenges and engage the larger community in these promising research directions.

3 Publications

3.1 Journal Papers

T. Lan, X. Lin, M. Chiang and R. Lee. How Bad Is Suboptimal Rate Allocation? In *IEEE/ACM Transactions on Networking*, 2011.

H. Mohsenian Rad, J. Huang, M. Chiang, and V. Wong, "Utility optimal random access: Reduced complexity, fast convergence, and robust performance", *IEEE Transactions on Wireless Communications*, vol. 8, no. 2, pp. 898-911, February 2009.

H. Mohsenian Rad, J. Huang, M. Chiang, and V. Wong, "Utility optimal random access without message passing", *IEEE Transactions on Wireless Communications*, vol. 8, no. 3, pp. 1073-1079, March 2009.

Jiayue He, Jennifer Rexford, and Mung Chiang, "Design for optimizability: Traffic management of a future Internet," to appear as a chapter in the book *Algorithms for Next Generation*

Architectures.

M. Yu, Y. Yi, J. Rexford, and M. Chiang, "Rethinking virtual network embedding: Support of path splitting and migration", *ACM Computer Communication Review*, April 2008.

3.2 Conference Papers

Yung Yi, Alexandre Proutiere, and Mung Chiang, "Complexity of wireless scheduling: Impact and tradeoffs", *Proc. ACM Mobihoc*, Hong Kong, China, May 2008.

Yung Yi, Junshan Zhang, and Mung Chiang, "Effective throughput and delay in wireless scheduling: Vacation model for complexity", *Proc. ACM Mobihoc*, New Orleans, LA, May 2009.

Wenjie Jiang, Rui Zhang-Shen, Jennifer Rexford, and Mung Chiang, "Cooperative content distribution and traffic engineering in an ISP network," in *Proc. ACM SIGMETRICS*, June 2009.

Jiayue He, Rui Zhang-Shen, Ying Li, Cheng-Yen Lee, Jennifer Rexford, and Mung Chiang, "DaVinci: Dynamically Adaptive Virtual Networks for a Customized Internet," in *Proc. ACM SIGCOMM CoNext Conference*, December 2008.

S. Liu, R. Zhang-Shen, W. Jiang, J. Rexford, and M. Chiang, "Performance bounds for peer-assisted live streaming", *Proc. ACM Sigmetrics*, Annapolis, MD, June 2008.

D. Xu, M. Chiang, and J. Rexford, "Link-state routing with hop-by-hop forwarding can achieve optimal traffic engineering", *Proc. IEEE INFOCOM*, Phoenix, AZ, April 2008.

3.3 Invited Presentations

Institute of Pure and Applied Mathematics Workshop on Network of Networks, "Content-Pipe Divide", November 2008.

Institute of Pure and Applied Mathematics Workshop on Mathematical Frontiers of Networking Research, "Distributed Scheduling", November 2008.

UCLA EE Department Seminar, "Two Open Problems in Networking: Random Access Performance and P2P Streaming Capacity", December 2008.

University of Toronto Networking Research Seminar, "Wireless Scheduling", April 2009.
Microsoft Research Seminar, "Wireless Scheduling", June 2009.

KAIST Information Science Seminar, “The Content-Pipe Divide”, June 2009.

IEEE WiOpt Plenary Speech, “Wireless Scheduling”, June 2009.

IEEE ICME Workshop on Emerging Technologies for Multimedia Communications Keynote Speech, “The Content-Pipe Divide”, June 2009.

MIT LIDS Student Conference, An Axiomatic Theory of Fairness, January 2010.

IEEE INFOCOM Plenary Panel, What is Good Research in Network Theory, March 2010.

K. U. Leuven 15th Simon Stevin Lecture on Optimization and Engineering, Optimization in Networking, July 2010.

Telcordia Annual Strategic Research Review Keynote Speaker, Optimization in Networking, July 2010.

10th International Symposium on Modeling and Optimization (MOPTA) Plenary Speech, ”Optimization in Networking”, August 2010.

Stanford ISL Colloquium, An Axiomatic Theory of Fairness, October 2010.

Caltech IST Lunch Bunch Talk, An Axiomatic Theory of Fairness, October 2010.

Caltech EE Seminar, Can Random Access Be Optimal?, October 2010.

U. Delaware EE Seminar, Can Random Access Be Optimal?, November 2010.

IPAM Optimization and Engineering Workshop, Optimizers on WiFi-Drivers, December 2010.

AFOSR Complex Network Workshop, Small World Delay, December 2010.

IEEE GLOBECOM Plenary Panel, Content Pipe Divide, December 2010.

4 Contributions

4.1 To the Principal Discipline

The 3D tradeoff analysis provides a systematic an fair comparison of a variety of wireless scheduling algorithms.

The adaptive CSMA algorithm is the first utility optimal random access algorithm.

The primal decomposition view enables a rigorous study on how to make network virtualization adaptive to dynamic network conditions and user requests without sacrificing efficiency or stability.

The computation of streaming capacity establishes benchmarks for all P2P streaming protocols to compare against.

The characterization of CDN-ISP interactions represents another step towards a systematic understanding of the interactions between those that operate the network and those that distribute content over the network.

4.2 To Other Disciplines

New mathematical methods in stochastic approximation theory were developed in the process of proving utility optimal of the adaptive CSMA algorithm.

Applications of economic theory to CDN-ISP interactions raise new questions on distributed versions of efficient market mechanisms.

The combinatorial algorithms developed as part of the solution to compute P2P streaming capacity solve novel problems in graph theory, such as minimum price tree with degree bounds.

4.3 To Development of Human Resources

At Princeton University, the project has involved a female student and a female post-doctoral researcher. In addition, the new research results have been incorporated in graduate-level courses in network optimization (EE) and data networking (CS), as well as special lectures in the undergraduate optimization and networking courses.